

SGIP PROGRAM PORTAL

Intent, Observations, and Technical Limitations

March 21, 2016

Presented by:

Tim O'Keefe

Director

Energy Solutions

Overview

- Who we are (Energy Solutions)
- Critical Design Requirements of SGIP Portal
- Summary of Program Open Results
- Application Submission Process
- Concerns from Program Open
- Technical Limitations and “First-Come First-Served”



Energy Solutions

- Energy Solutions is a DSM consultancy founded 1995
 - Approximately 25% of work is Information Systems related
- Information systems work includes:
 - Design, develop, and maintain online portals for many programs in EE, DG, and DR
 - Our online systems have processed millions of incentives totaling billions of dollars



Energy Solutions

- In 2011, we developed the SGIP online portal for PAs
 - Migrated program from individual PA MS Access DBs
 - Original portal for use by PAs only - applications were submitted by mail and email
- Throughout 2015 we designed and re-developed the SGIP database portal to include an Applicant Interface



Design Requirements and Intent

- Primary design requirements for the portal:
 - **Usability** for applicants and program administrators
 - **Workflow** management of incentive application process
 - **Management** of documents, communications, payments, and program budgets
 - **Calculation** of incentives
- Expected annual activity: 1,000 applications statewide
 - Approximately 2x historic average
- Designed to be highly scalable if program grows due to increased budget or higher volume of smaller incentives (Residential Storage, for example)



Summary of Submissions

- On February 23 at 8:00am, the SGIP opened with ~\$44.5M in available incentives
 - ~\$32.5M "Level 2" (Emerging Tech and Renewable)
 - ~\$12M "Level 3" (Non-Renewable)
- Demand was extremely high
 - Level 2 requested incentives were over 6x available budget

	Level 2		Level 3	
	Apps	Incentives	Apps	Incentives
<i>Budget</i>		\$32,454,041		\$12,084,536
Total	901	\$206,825,739	9	\$2,477,240
First 10 mins	655	\$179,268,913	6	\$2,195,000



Application Submission

Submission Workflows

- Most applicants completed and verified applications prior to program open
 - Only had to open the application and click “Submit” when program opened
- For these applicants, there were two primary workflows to submit using the user interface:
 - “Complete” workflow
 - “Bookmarked” workflow



“Complete” Workflow

Workflow Step (Click)	Database Load
<i>Log in</i>	<i>Light</i>
View list of applications	Heavy
Open Edit Application view	Heavy
Load the Submit Application page	None
Agree to the Terms of Use	None
Click the Submit button	Light

- 6 clicks from accessing the system to submission
 - Login step unnecessary from browsers used to access system recently
- Two heavy processes
 - Hundreds of database queries to load application data
 - Creating load on entire system
 - Highly performant under typical loading (~1 sec for each)



“Bookmarked” Workflow

Workflow Step (Click)	Database Load
<i>Log in</i>	<i>Light</i>
Load the Submit Application page (from Bookmark)	Light
Agree to the Terms of Use	None
Click the Submit button	Light

- 3 clicks from accessing the system to submission
- Login step unnecessary for browsers used to access system recently
- No heavy processes



Other Submissions

- Some Applicants attempted to send the Submit command directly without using the user interface
 - Submit command hidden until 8:00am program open
 - Present in other parts of site
 - Applicants succeeding with this technique did so with a combination of research and guesswork
 - Some Applicants attempted this and ‘guessed’ incorrectly
 - Most Applicants attempting this technique also had people using the “Complete” or “Bookmarked” workflows as backup
 - This method was very light on system resources, so responses returned quickly and the system was minimally impacted



Concerns from Program Opening

Concerns

Industry stakeholders have raised several concerns with the performance of the system during the program opening:

- Extremely slow performance
- "Error" messages preventing login
- Users unable to login after multiple attempts
- Duplicate application IDs were issued
- Timestamps not in the order submitted



Concerns

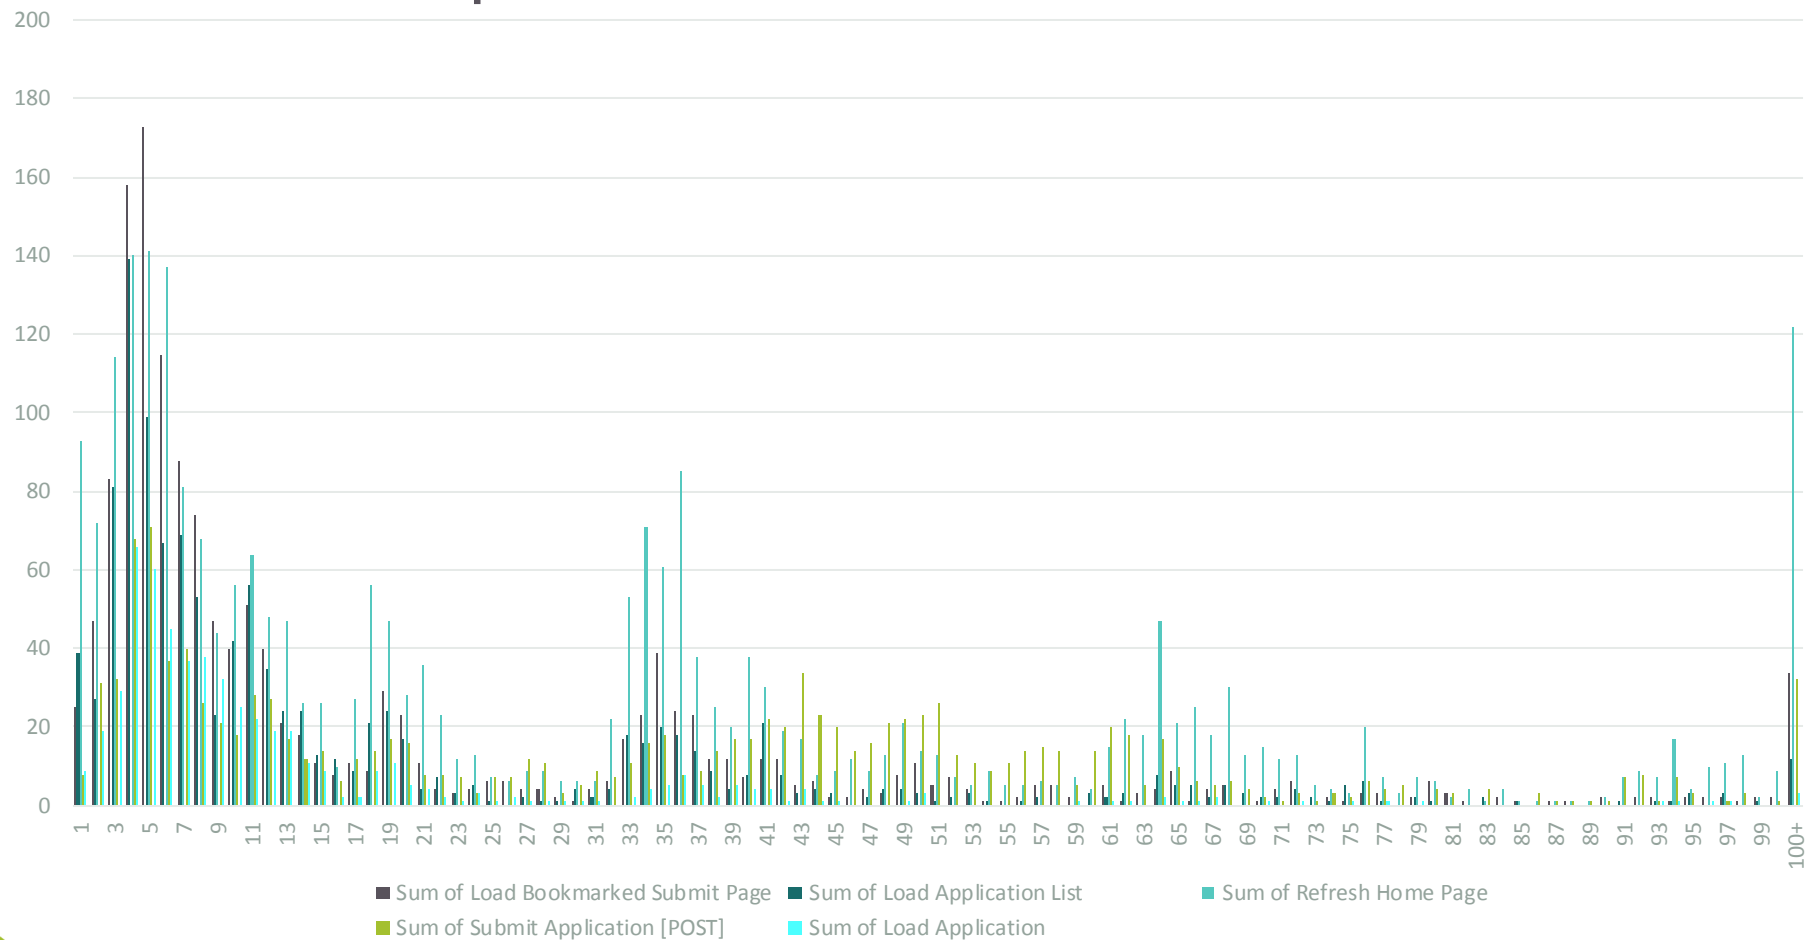
Industry stakeholders have raised several concerns with the performance of the system during the program opening:

- Extremely slow performance
- "Error" messages preventing login
- Users unable to login after multiple attempts
- Duplicate application IDs were issued
- Timestamps not in the order submitted



Slow System Performance

System performed very slowly between 8:00 and 8:10am
~20% of requests took between 30 and 75 seconds



Causes of Slow Performance

Unusually high demand

- ~28,000 server requests in the first 10 minutes
 - Approximately 50 per second = 1 per 20 milliseconds

Significant redundant activity

- Some Applicants had many registered users
- Some Applicants had many users sharing a single login
- Half of Applicants attempted* duplicate submissions
 - 2,200 submission attempts for 650 applications in first 10 minutes
 - Average of 3.4 attempted submissions per application program-wide
- This activity slowed the system for all users

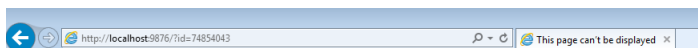


*Note that the system does not allow the creation of duplicates and none were 'created' by these attempts

Effects of Slow Performance

Client Timeouts

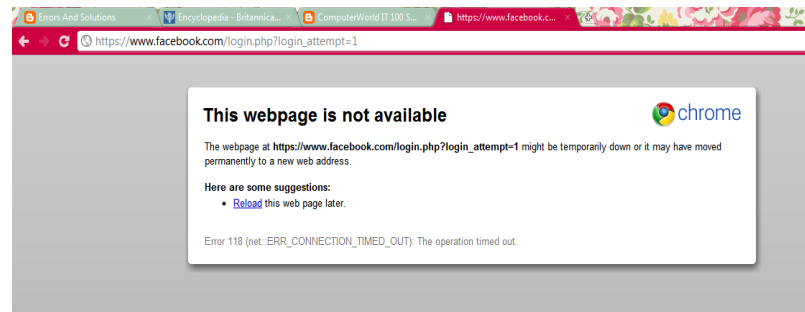
- Web browsers typically “time out” after waiting more than a certain time
 - “Time out” means your browser will wait for a response, and then give up* after a specified time
 - Varies among browsers from 1 – 5 minutes (or user can set)
- Timeouts are the cause of:
 - "Error" messages preventing login
 - Users unable to login after multiple attempts



This page can't be displayed

- Make sure the web address http://localhost:9876 is correct.
- Look for the page with your search engine.
- Refresh the page in a few minutes.

Fix connection problems



Effects of Slow Performance

Duplicate Application Codes

- When a submission is processed, a sequential application code is assigned
- Assignment requires database lookup
 - “What is currently the highest app code?”
- When the server is under unusual load, multiple lookups were open concurrently
 - Application waits to hear what highest app code is before writing one
 - If many are waiting at the same time, they may all get the same response and write the same next app code

NOTE: This has no effect on timestamps, which do not require a database lookup



System Performance Summary

1. Timeouts were the primary errors seen by users, and prevented some users from accessing the system
2. A '500 error' was returned for 15 duplicate submittal attempts (projects had already received timestamp)
3. System security was not compromised, and customer data was protected and secure at all times
4. The system was not 'hacked' or 'attacked' – all access was authorized and intended



Technical Limitations

The Internet and “First-come, First-served”

Timestamp Ordering

Concern:

4. Perhaps most importantly – some parties are reporting that the order in which they submitted their applications does not match the time stamp order on the released SGIP data. This mismatch is significant in that it indicates that the data published on March 1, 2016 is either factually incorrect or, equally problematic, the portal was so over-loaded with the number of applications at 8:00 a.m. that the time stamps were potentially allocated incorrectly to individual applications.



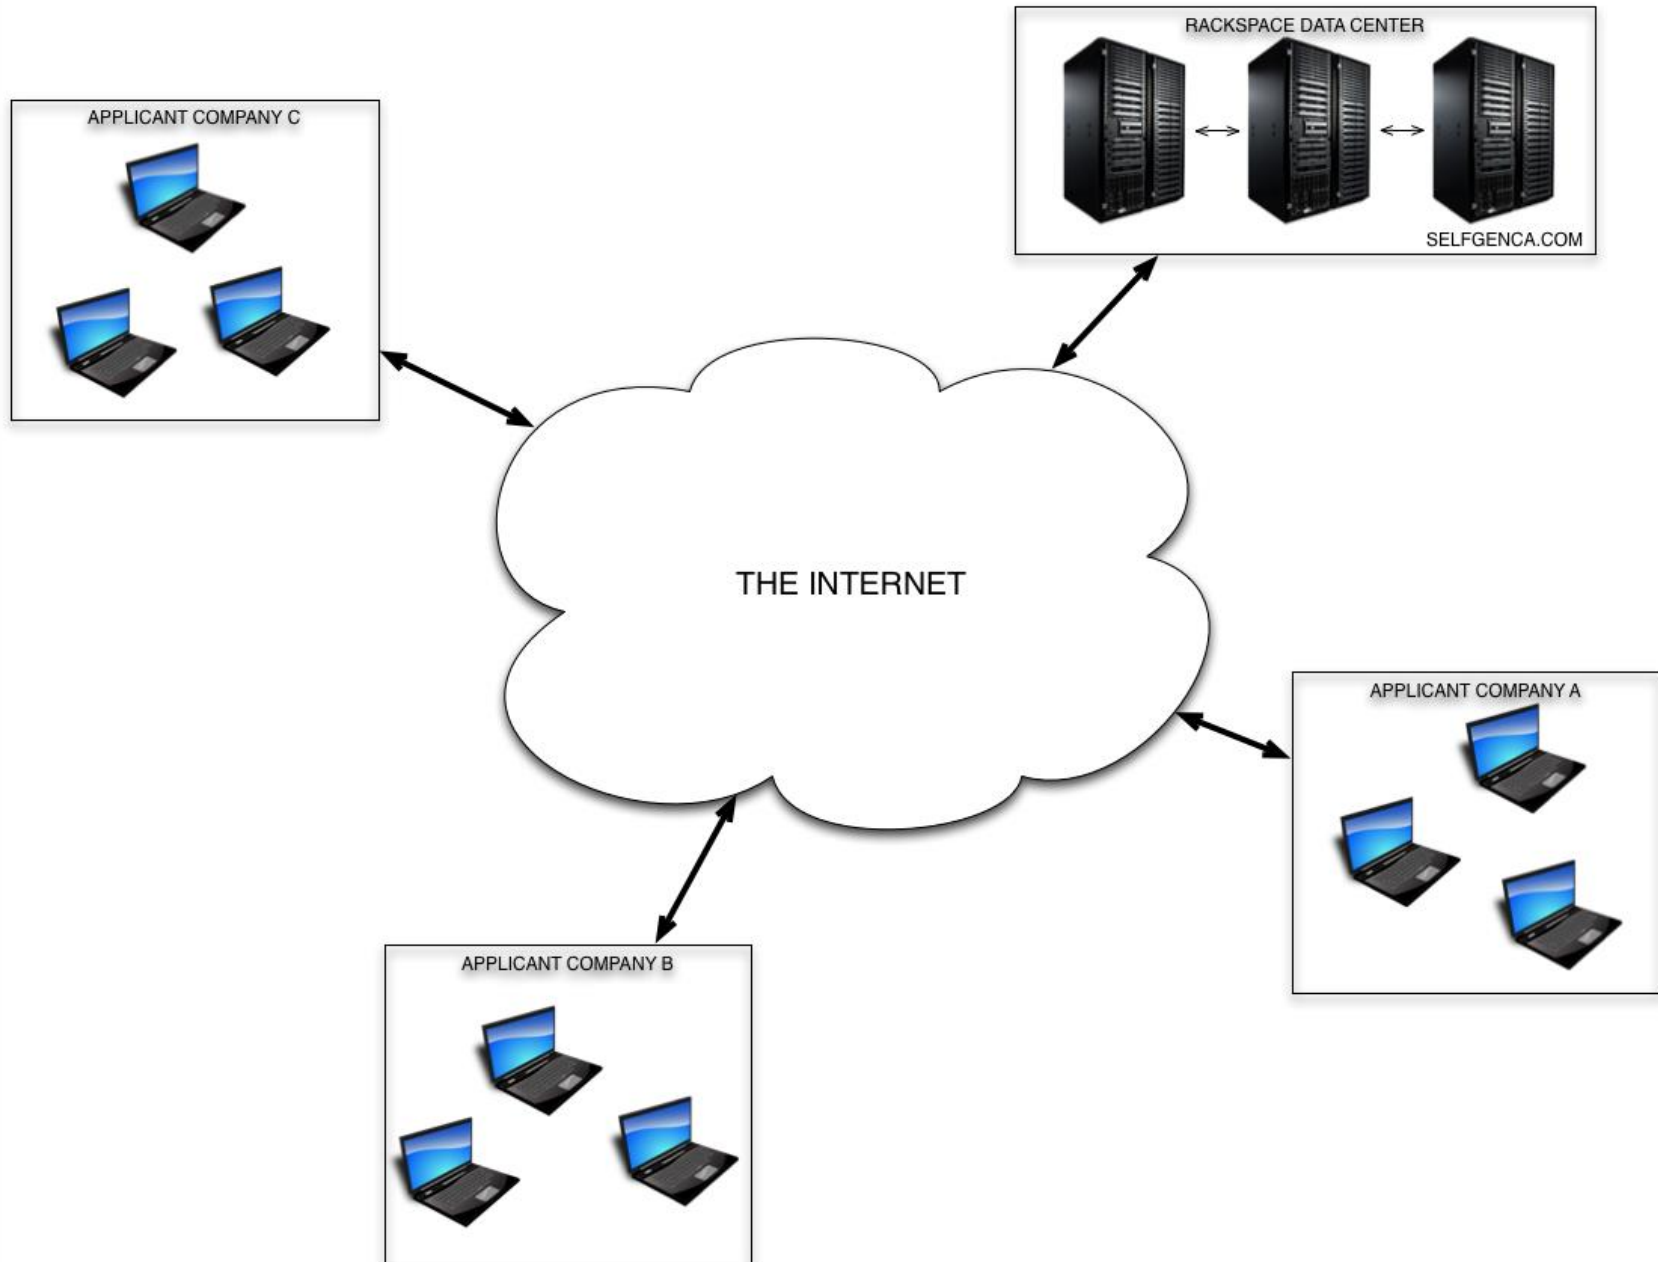
Timestamp Ordering

Concern:

4. Perhaps most importantly – some parties are reporting that the order in which they submitted their applications does not match the time stamp order on the released SGIP data. This mismatch is significant in that it indicates that the data published on March 1, 2016 is either factually incorrect or, equally problematic, the portal was so over-loaded with the number of applications at 8:00 a.m. that the time stamps were potentially allocated incorrectly to individual applications.



Timestamp Ordering



Timestamp Ordering

World Internet Topology

Brought to you by **AT&T Labs**

Powered by **GLUMETA**

This map represents the backbone of the Internet as of August 2005. Each line depicts the physical engineering route from a host computer to each of more than 125,000 network nodes around the world. The map does not represent the physical or logical location of servers, but rather is a topological representation of the complex network that enables us to use the Internet. It shows the possible number of interconnected networks owned and managed by private companies. These networks combine to form the world Internet backbone in individual. This map is a product of the Internet Mapping Project. It was compiled and created by Bill Chesney and Steven Bello at AT&T Labs Research, using technology and methods developed by the Luminet Corporation.

AT&T's Network by the Numbers.

9.81

Petabytes of data transmitted across AT&T's network on an average business day. It is the equivalent of moving the entire electronic contents of the Library of Congress every three minutes.

1

AT&T's rank among broadband providers in the United States.

12.9 Million

AT&T broadband customers in America.

540,000

Miles of Internet backbone fiber AT&T owns and operates.

\$6 Billion

Amount AT&T will spend by 2008 to bring fiber optics deeper into neighborhoods.

36

AT&T Internet data centers around the world.

301,760

AT&T employees worldwide.

97%

Percentage of the world economy reached by AT&T's networks.

99.998%

AT&T network reliability.

49,000

Number of World Headquarters AT&T provides or enables.

166

Number of United States cities where AT&T offers 3G wireless high-speed Internet access.

3 Million

AT&T network's average data subscription.

160%

Increase in bandwidth demand per AT&T user between July 2004 and October 2005.

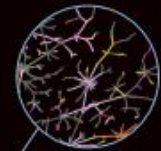
7

Patent Peaks awarded for inventions developed at the AT&T Bell Labs.

2

Average number of new patent applications AT&T files - every day.

Legend: AT&T, Bell Labs, Sprint, Verizon, Time Warner, Comcast, Charter, Cox, Earthlink, NetScout, Level 3, Qwest, CenturyLink, Windstream, SBC, AT&T Mobility, AT&T Knowledge Ventures, AT&T World of Outdoors, AT&T Worldnet, AT&T Worldnet TV, AT&T Worldnet TV HD, AT&T Worldnet TV HD+.



These clusters represent the most densely populated network hubs. Each cluster represents a unique provider capable of connecting thousands of individual users to the Internet. The complexity of the network structure can make it difficult to distinguish the major components of the network, but the overall structure is clear.

World Internet by the Numbers.

More than

320,000

Individual network nodes listed by the Internet Mapping Project.

48 Million

Users on the Internet in 2005. (Source: IDC)

1.133 Billion

Internet users in 2006. (Source: Internet World Watch)

6.4 Million

New Internet users getting online every month. (Source: Internet World Watch)

1.6 Billion

Internet users in use in 2006. (Source: IDC)

40 Million

New DNS hosts every year. (Source: Internet World Watch)

35,000

New pages added to the Internet every day. (Source: Internet World Watch)

100 Million

YouTube videos downloaded every day. (Source: YouTube)

161

Exabytes of new electronic data created every year. (Source: IDC)

12 Million

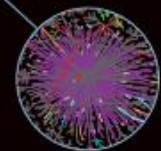
Miles of new fiber deployed in 2005. (Source: Telecommunications Industry Association)

15 Million

Miles of new fiber to be deployed globally by 2006. (Source: Telecommunications Industry Association)

\$72.5 Billion

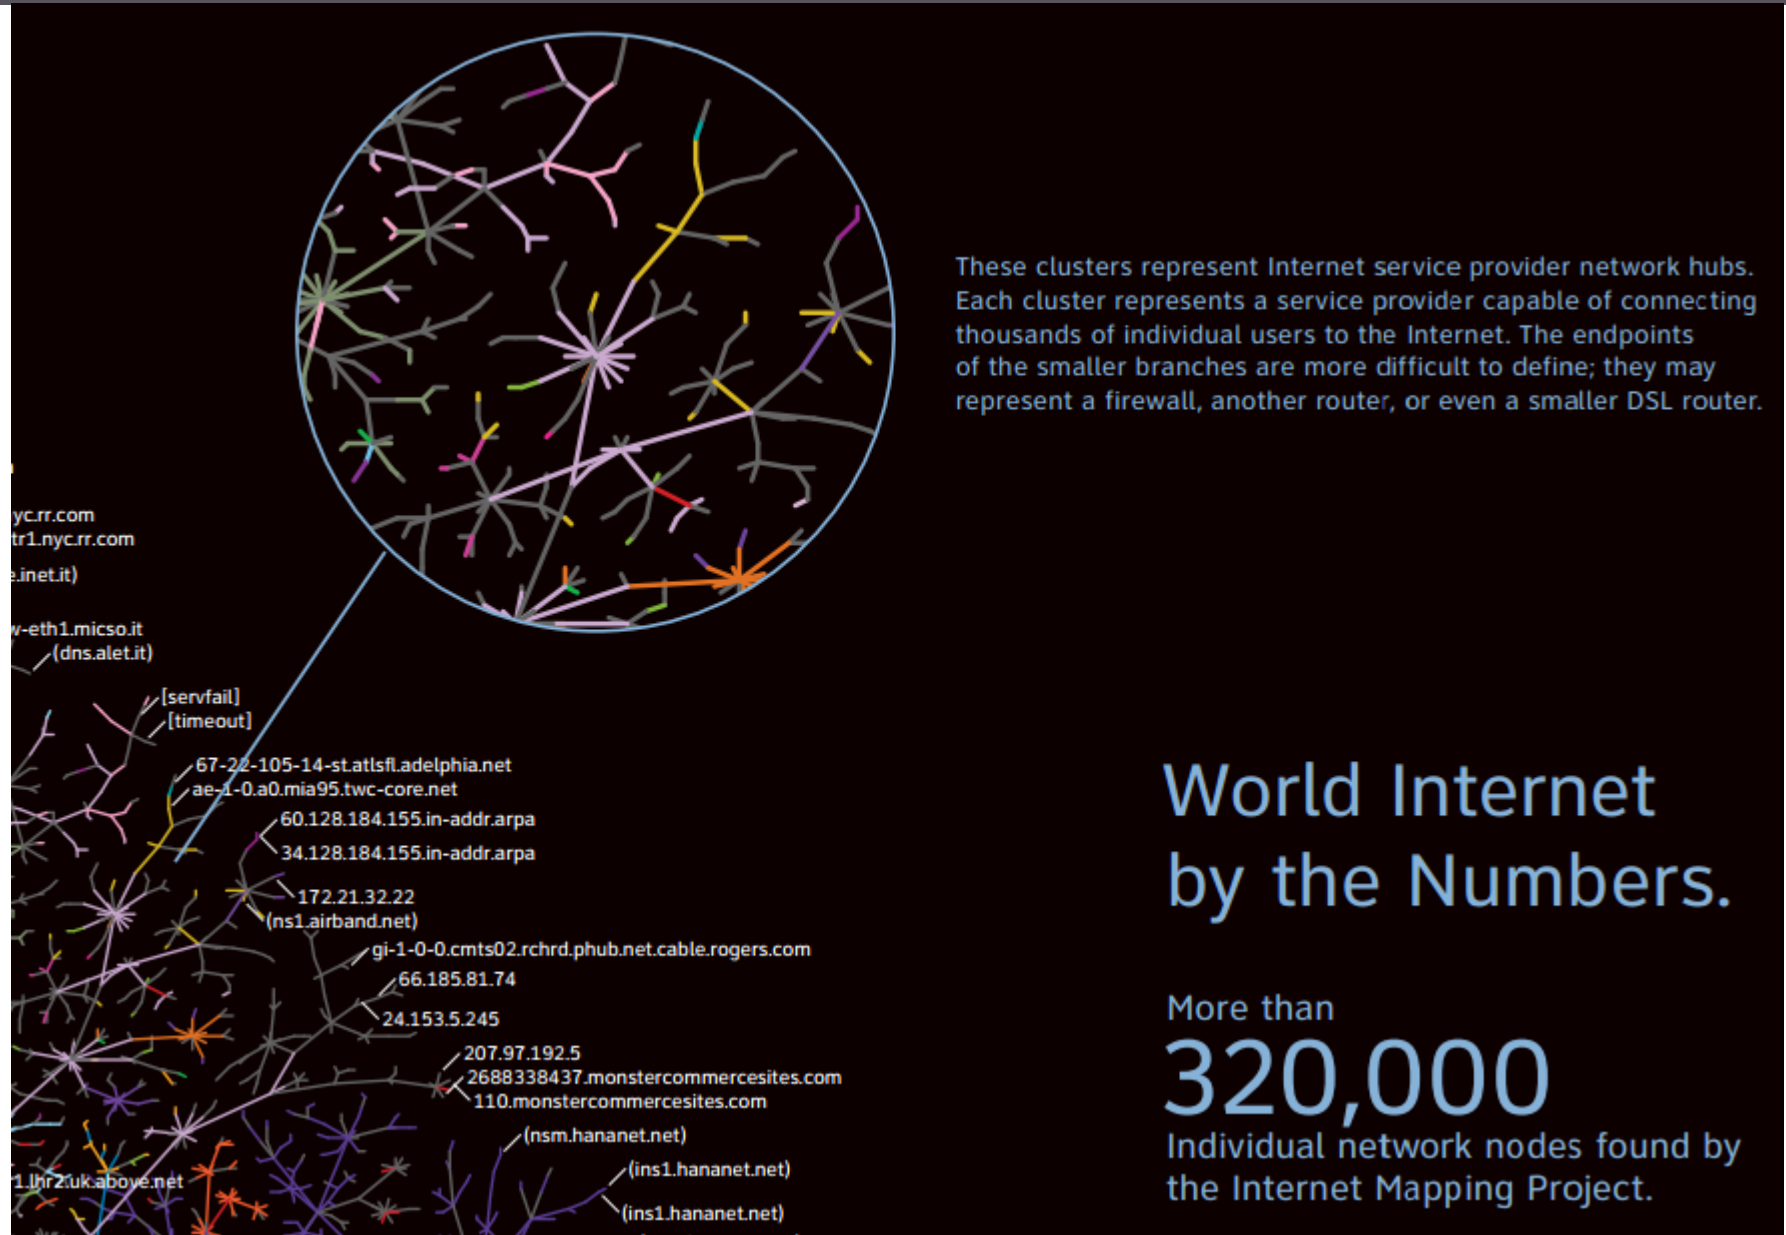
Annual spending on research and development in the United States by 2006. (Source: Telecommunications Industry Association)



These large groups represent major hubs. They are a mix of both of global connectivity and local connectivity. The complexity of the network structure can make it difficult to distinguish the major components of the network, but the overall structure is clear.



Timestamp Ordering



Timestamp Ordering

Internet Latency

Each request sent across the internet is routed many times across various networks

- Typically 10-20 network 'hops' occur for each request
- Each 'hop' can have congestion and variable latency from 10 – 1000+ milliseconds
- The path of 'hops' may be different for each request, even from the same computer to the same server
- The server has no information about when the request left the requestor – it only knows when it arrived



Timestamp Ordering

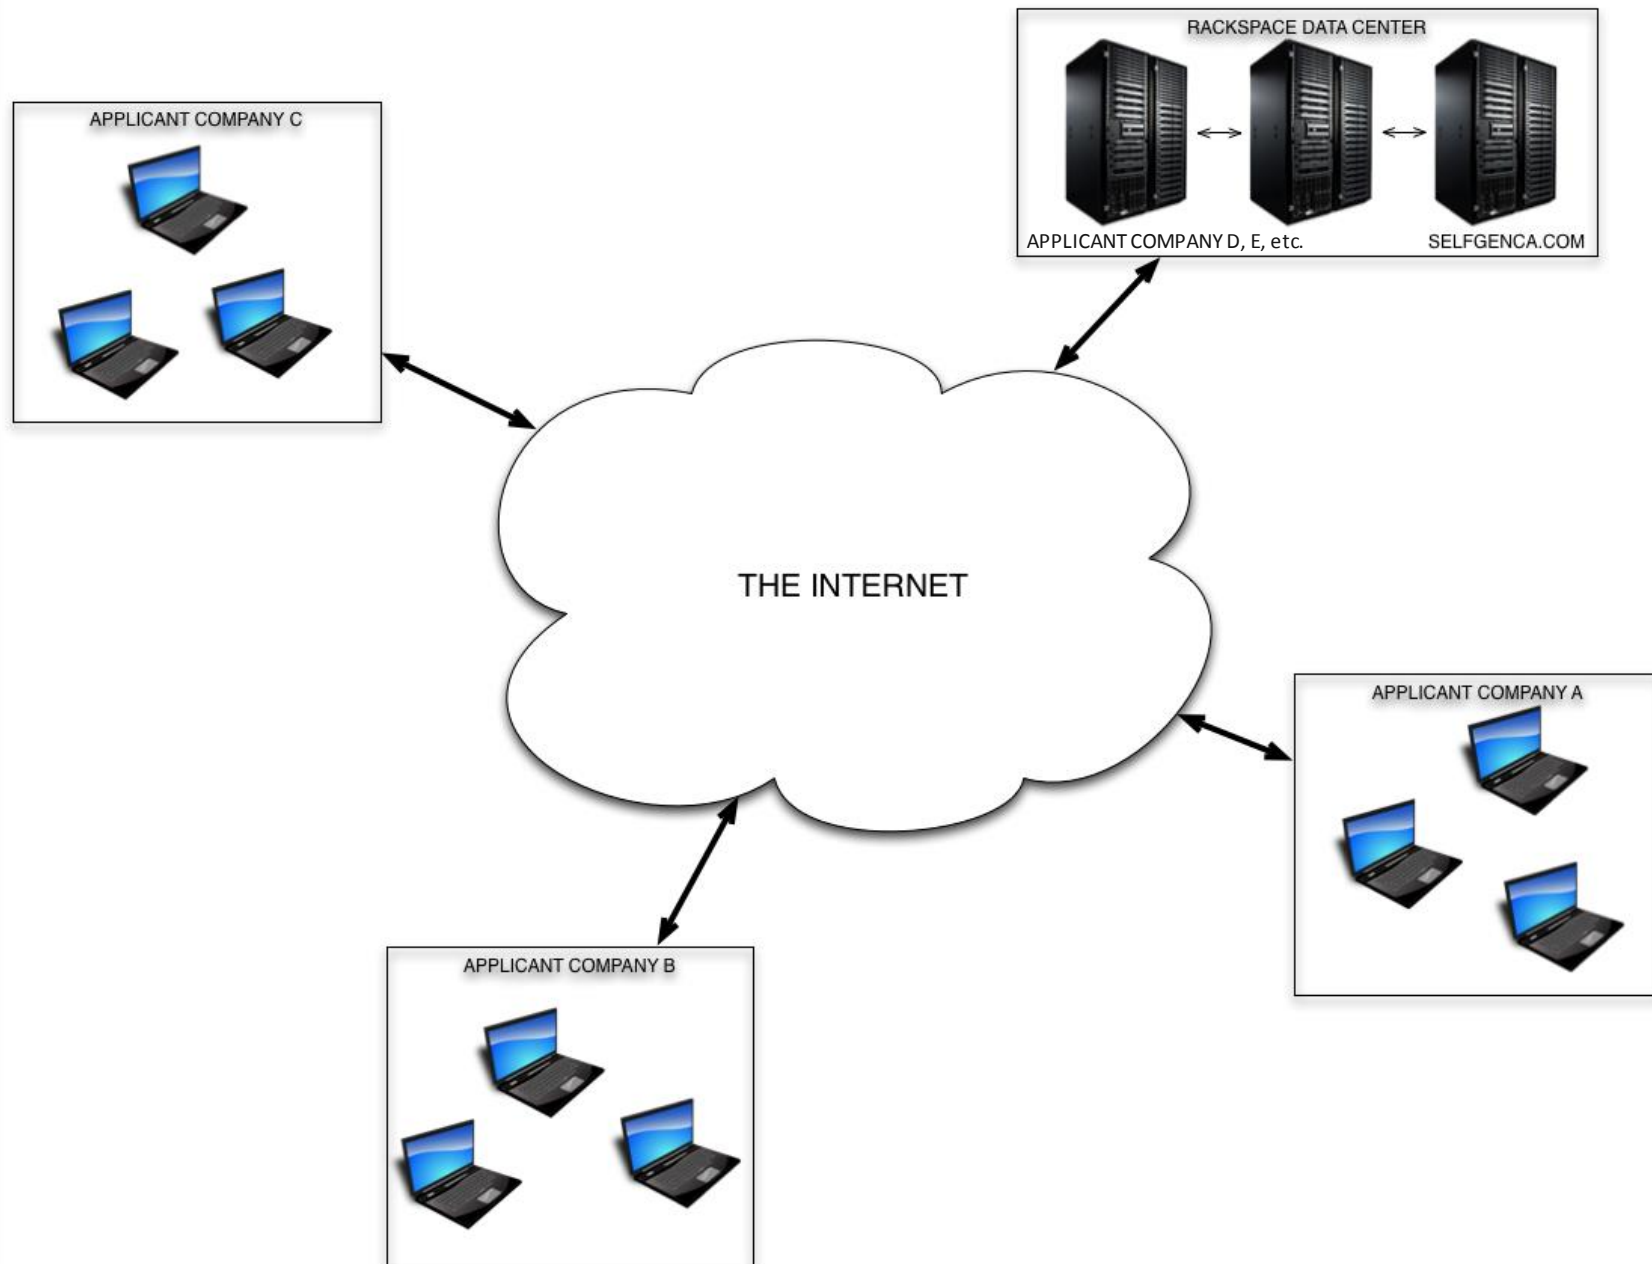
Internet Latency

The result: The order of clicks has no bearing on the order requests are received at the server (at least below ~10 second granularity)

* Note that some applicants used computers within the Rackspace physical network where the SGIP online portal is hosted to submit their applications, eliminating the effects of internet latency for their submissions



Timestamp Ordering



Timestamp Ordering

Server Congestion

Even at the server level, we know that requests do not necessarily complete in the order received (at a millisecond level) due to system “multi-threading” and various internal system queueing technologies

- This is particularly true under heavy load and when processes have different resource requirements
- Similar to the line at a grocery store
 - You may end up in line behind someone with many items



Timestamp Ordering

Timestamp Assignment

Timestamps are assigned by necessity as the very last step of the process

- Can't assign timestamp unless process completes successfully
- Timestamp designates when the submission was 'received'



Conclusion

Considering the volume of applications

- 650* in 10 minutes = ~1 application / second

and the inherent latency of network technology

- Requests can take +/- 5 seconds in unloaded conditions, more variance under load

Internet submission is not capable of ordering applications in the order 'submitted' at the browser with this concentration of demand, even with no server slowdown



*650 *successful* submissions – all in all, there were 2,200 submission attempts in first 10 minutes